Network-wide Packet Behavior Identification and Private Network Function Outsourcing

Huazhe Wang
University of Kentucky
huazhe.wang@uky.edu

Chen Qian
University of Kentucky
qian@cs.uky.edu

arXiv:1603.02613v1 [cs.NI] 8 Mar 2016

ABSTRACT

Identifying the network-wide forwarding behaviors of a packet is essential for many network management applications. We present AP Classifier, a control plane tool for packet behavior identification. Experiments show that the processing speed of AP Classifier is faster than existing tools by at least an order of magnitude. Furthermore, AP Classifier uses very small memory and is able to support real-time updates. We also present a network function outsource framework with AP Classifier which can provide security properties.

Keywords

Packet behavior identification; Software defined networking; Network function

1. INTRODUCTION

Network-wide packet behavior identification is a function that discovers the actual behaviors of the packets including their forwarding path, where they stop, and which boxes they transverse. It is essential for many network management applications, including rule verification, policy enforcement, attack detection, traffic engineering, and fault localization. Current tools [3][4][6] that can perform packet behavior identification either incur large time and memory costs or do not support real-time updates. We have designed and implemented AP Classifier, a software defined networking (SDN) control plane tool for packet behavior identification. Experiments show that the processing speed of AP Classifier is faster than existing tools by at least an order of magnitude. Further, AP Classifier uses very small memory and is able to support real-time updates.

Due to the high complexity and cost of managing network functions (also referred to as middleboxes), [5] have explored the possibility for enterprises to outsource the processing of their traffic to third-party clouds. To adopt this innovation, an enterprise has to provide the detailed configurations of these network functions which may leak sensitive policy rules to potential attackers. Facing the challenge, we propose to design a privacy-preserving framework for private network function outsourcing using AP Classifier.

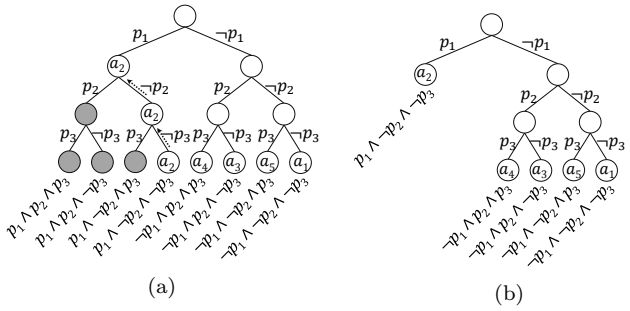

Figure 1: A sample AP Tree

2. AP CLASSIFIER

AP stands for atomic predicates, a concept for a network developed in [6]. The packets that are evaluated to true by the same atomic predicate have identical behaviors at all boxes. With the algorithms in [6], we calculate the set of predicates and atomic predicates of a network. AP Classifier performs two-stage processing for a packet. First, using the AP Tree, it classifies the packet to the atomic predicate that evaluates to true for the packet. Second, AP Classifier determines all behaviors by using the atomic predicate, network information, and ingress box of the packet.

2.1 AP Tree

AP Tree is a novel binary tree structure constructed by labeling nodes on each level with a predicate of the network. Let $P = \{p_1, p_2, ..., p_k\}$ be the set of predicates of the network. The root is labeled by p_1. At level i, the 2^i internal nodes are each labeled by p_i. Starting from the root, at each internal node, the input packet is evaluated by the predicate in the label. If the result is true, the packet continues to be evaluated in the left sub-tree. Otherwise it goes to the right sub-tree. A leaf node is then labeled by $q_1 \wedge q_2 \wedge ... \wedge q_k, q_i \in \{p_i, \neg p_i\}$, which specifies the set of packets reaching the leaf. From the definition of atomic predicates, leaf labels (that are not false) represent the atomic predicates of P. Fig. 1(a) shows a sample AP Tree of three predicates. Subtrees that specify an empty set can be pruned since no packet can reach their leaves. Assuming shaded nodes in Fig. 1(a) are empty, Fig. 1(b) shows the AP

Tree after pruning.

AP Tree optimization To increase query throughput of the AP Tree, we have designed a heuristic algorithm to construct an AP Tree with minimum average leaf depth. The intuition is that we can obtain different AP Tree patterns and average leaf depths if we label nodes with predicates in different orders. A pair-wise relation between predicates is developed which enable AP Classifier to determine which predicate to select at each internal node easily. Data plane state of two real networks [2][1] are used to construct AP Trees. Average depth of AP Trees constructed using AP Classifier is 10.6 and 16.8 respectively, which are at least 50% smaller than labeling predicates in random orders.

Dynamic updates Network dynamics, including link and rule changes, can be represented as addition and deletion of predicates. New added predicates are placed at the bottom of the AP Tree and form new leaves. We do not change the AP Tree if some predicates are deleted since it still ensure correctness of classification. After a large number of updates, an AP Tree structure is no longer optimum. Hence, AP Classifier reconstructs the AP Tree on a second process. During reconstruction, the original process still maintains the old AP Tree by performing updates, and responding to queries.

2.2 Computing Packet Behaviors

Since the atomic predicate is in the form $q_1 \wedge q_2 \wedge ... \wedge q_k, q_i \in \{p_i, \neg p_i\}$, for any predicate p_j, AP Classifier can easily check whether the predicate evaluates to true or false for the packet. Recall that p_j represents a packet filter of an ACL or output port. Hence AP Classifier can determine whether the packet is dropped and which port it is forwarded to at each box.

2.3 Some Experimental Results

We evaluate the performance of AP Classifier using the data plane network state from [2] and [1], including forwarding tables and access control lists (ACLs). Our results show that AP Classifier, running on a general purpose desktop computer, uses a few MBs memory and supports more than two millions of queries per second. In addition it can be updated in real time (¡4 ms for 95% updates).

3. A PRIVATE NETWORK FUNCTION OUTSOURCING FRAMEWORK

Fig. 2 shows a privacy-preserving framework for network function outsourcing using AP Classifier. A local agent A_1 at the enterprise tunnels both ingress and egress traffic to the cloud. Traffic are classified to a label at the local agent, sent to the cloud with the label and sent back to the local agent after in-cloud processing. Before traffic entering the cloud, packet head fields are encrypted if needed. All processing and forwarding

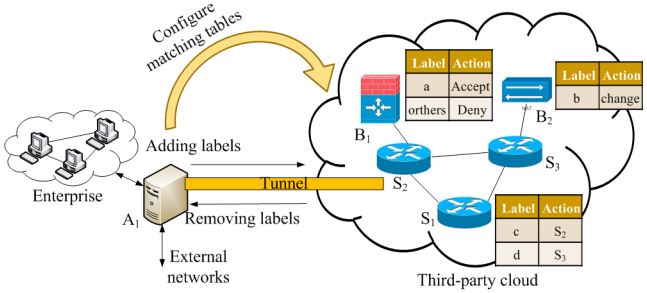

Figure 2: A processing chain outsource framework.

behaviors on middleboxes and switches in the cloud are based on labels instead of rules. For example, switch S_1 forwards packets with a label c to switch S_2 and packets with a label d to S_3. Similarly, firewall B_1 only accepts packets with a label a and drops other packets. Under this mechanism, processing policies and packet headers are out of sight of third-party cloud.

To achieve the framework described above, we calculate atomic predicates using all rules on switches and middleboxes. The set of atomic predicates are mapped to a set of labels. At the local agent, we classify each packet to an atomic predicate using AP Classifier before they enter the cloud. A predicate of the network which is a disjunction of a subset of atomic predicates can be mapped to a set of labels. Rules on switches and middleboxes are firstly converted to predicates and then mapped to labels. For example, an ACL list can be presented as a predicate A, then A is mapped to a set of labels $\{a_1, a_2, ..a_k\}$ corresponding to atomic predicates $\{p_1, p_2, ...p_k\}$ whose disjunction is A.

Some middleboxes may change packet headers, we modify the labels of the packets instead in our design. Considering a box which changes packet headers from h_1 to h_2, the atomic predicate that h_2 belongs to is calculated proactively, denoted as p. The action at the box is configured as changing labels of the matched packets to the labels corresponding to p. Then the packets continue rest of processing using the new labels.

4. CONCLUSION AND FUTURE WORK

We have proposed AP Classifier for network-wide packet behavior identification that can process millions of queries per second. It uses only a few MBs memory and is robust under dynamic data plane changes. A framework for private network function outsourcing using AP Classifier is also presented.

5. REFERENCES

[1] Header space library and netplumber. `http://bitbucket.org/peymank/hassel-public/`.

[2] The internet2 observatory data collections. `http://www.internet2.edu/observatory/archive/data-collections.html`.

[3] T. Inoue, T. Mano, K. Mizutani, S. Minato, and O. Akashi. Rethinking packet classification for global network view of software-defined networking. In *Proc. of IEEE ICNP*, 2014.

[4] P. Kazemian, G. Varghese, and N. McKeown. Header space analysis: Static checking for networks. In *Proc. of USENIX NSDI*, 2012.

[5] J. Sherry, S. Hasan, C. Scott, A. Krishnamurthy, S. Ratnasamy, and V. Sekar. Making middleboxes someone else's problem: network processing as a cloud service. *ACM SIGCOMM Computer Communication Review*, 42(4):13–24, 2012.

[6] H. Yang and S. S. Lam. Real-time verification of network properties using atomic predicates. In *Proc. of IEEE ICNP*, 2013, extended version in *IEEE/ACM Transactions on Networking*.